Publication as Prestige Theater: AI Startups, Institutional Science, and the Politics of Epistemic Legitimacy

Douglas C. Youvan

doug@youvan.com www.youvan.ai

July 29, 2026

A recent analysis of publication activity among artificial-intelligence startups has been presented as evidence of a troubling paradox: companies claiming to transform science rarely participate in conventional scientific publishing. The finding is descriptively interesting but has been rhetorically enlarged beyond what the underlying data can support. Publication frequency is treated as a proxy for scientific seriousness, transparency, reproducibility, and social responsibility, although these properties are neither identical nor reliably produced by academic publication itself. The resulting argument functions partly as institutional reassurance. At a moment when corporations possess greater computing resources, larger engineering teams, and greater public influence than many universities, publication metrics restore academic scientists to the position of epistemic gatekeepers. This paper examines the political and institutional assumptions embedded in that framing. It argues that the real issue is not whether AI startups publish enough papers, but whether their consequential claims can be independently tested, reconstructed, challenged, and corrected. Scientific accountability may be served through papers, but also through code, data, model weights, technical documentation, external audits, reproducible evaluations, and public disclosure of failures. A publication-centered critique therefore risks defending the prestige economy of institutional science rather than establishing a serious standard for knowledge produced outside it.

Keywords: artificial intelligence, scientific publishing, epistemic legitimacy, institutional science, reproducibility, AI startups, academic prestige, open science, scientific accountability, bibliometrics.

Google Drive: https://drive.google.com/drive/folders/15Kd8bUjm7zVaKdVIwbPhqjjV68aUZQvc

A collaboration with OpenAI GPT-5.6 Thinking. CC4.0.

https://orcid.org/0009-0004-4206-4253

1. The Comforting Scandal

The claim is immediately satisfying to institutional science: artificial-intelligence companies are valued in the billions, promise to transform medicine and discovery, and yet publish remarkably few scientific papers. The conclusion appears obvious. These firms possess money, computing power, and cultural influence, but they do not participate adequately in the institutional practices by which legitimate knowledge is established. Academia, by contrast, still publishes, reviews, cites, archives, and debates. Corporate AI may be wealthy, but institutional science remains epistemically superior.

This is an emotionally attractive story for researchers working in universities, government laboratories, nonprofit institutes, and scholarly publishing. It arrives during a period of institutional anxiety. Academic scientists increasingly depend on commercial cloud providers, proprietary models, privately owned datasets, corporate software platforms, and expensive computational infrastructure. The largest AI companies can recruit prominent researchers, absorb entire laboratories, and conduct experiments beyond the financial reach of most universities. Traditional institutions therefore face a widening gap between their inherited authority and their declining material control over the most advanced technical systems.

Publication offers a symbolic answer to this displacement. Universities may no longer control the largest models, the most powerful processors, or the most heavily capitalized research programs, but they still control the rituals by which scientific legitimacy is publicly recognized. A company may build a consequential system, but unless its work passes through familiar channels of authorship, peer review, citation, and archival publication, its accomplishments can be described as insufficiently scientific. The publication record thus becomes a defensive boundary around institutional authority.

The resulting controversy is not merely about whether AI companies communicate their findings. It is about who is permitted to define what counts as knowledge.

The article's headline transforms a narrow bibliometric observation into a broad indictment. AI startups are said to be "barely publishing their research." Yet the study does not directly determine how much research these firms conduct, how much they document internally, how much they disclose through technical reports, or how much outsiders can evaluate through released systems, code, benchmarks, model weights, patents, demonstrations, and independent testing. It counts a particular class of formal outputs and then allows those outputs to stand in for scientific participation as a whole.

That substitution is the article's central rhetorical maneuver. A count of indexed papers becomes a judgment about whether discoveries have been documented. A judgment about documentation becomes a judgment about validation. A judgment about validation becomes a

concern about reproducibility, energy use, safety, and social responsibility. Each transition sounds plausible, but none follows automatically from the one before it.

A company can publish many papers while disclosing little of practical value. It can omit training data, software dependencies, evaluation prompts, model weights, failure cases, safety incidents, energy consumption, and the undocumented interventions required to reproduce a result. The resulting literature may be formally public yet operationally opaque.

Conversely, a company may publish few journal articles while releasing code, weights, detailed technical reports, public benchmarks, or functional systems that can be examined by thousands of outside researchers. Such disclosure may still be incomplete, but it cannot simply be equated with the absence of scientific documentation.

The article therefore stages the wrong confrontation. The relevant distinction is not between companies that publish and companies that do not. It is between claims that can be independently examined and claims that remain protected from meaningful scrutiny.

2. Publication Is Not Science

Scientific publication is an important institution, but it is not identical to science. It is one mechanism by which scientific claims are communicated, evaluated, and preserved. At its best, publication makes evidence public, records methods, establishes priority, exposes arguments to criticism, and allows results to enter a cumulative body of knowledge. At its worst, it produces a ceremonial literature that satisfies professional requirements without making the underlying work transparent or reproducible.

The distinction matters because the article treats publication as though it carried an intrinsic epistemic authority. A company with few papers appears scientifically suspect, while an institution producing many papers appears to be fulfilling its scientific obligations. Yet publication volume alone reveals little about the reliability of the claims being published.

Academic science has long struggled with irreproducible findings, selective reporting, publication bias, statistical manipulation, inaccessible data, unavailable code, incomplete methods, honorary authorship, citation cartels, prestige effects, and the suppression of negative results. Peer review often evaluates whether a manuscript is plausible and professionally presented, not whether its central claims have been independently reproduced. Journals may certify that a paper has passed through an institutional process, but they rarely certify that its conclusions are true.

This does not render publication worthless. It shows only that publication is a fallible social technology rather than a sacrament.

The prestige of publication derives partly from its symbolic function. A paper places a claim within an authorized system. It gives the claim a title, an author list, an institutional affiliation, a journal name, a publication date, and eventually a citation count. These markers make knowledge administratively visible. They allow universities to rank researchers, granting agencies to distribute money, journals to calculate influence, and institutions to advertise productivity.

The system is therefore simultaneously epistemic and bureaucratic. It communicates knowledge, but it also allocates status.

That dual function is often obscured when scientists speak about publication as though it were simply the neutral record of discovery. A published paper is not merely a report of something learned. It is also a professional credential, a claim of priority, a unit of institutional output, a mechanism of career advancement, and an entry in a competitive prestige economy.

The article under discussion depends heavily on this prestige economy. Its indictment becomes persuasive only if the reader accepts that the formal paper is the normal and morally preferred unit of serious scientific contribution. Once that assumption is weakened, the headline loses much of its force.

The question then changes. Instead of asking why AI startups do not produce more papers, one must ask what forms of disclosure would allow their claims to be tested.

For a frontier model, the answer might include documentation of the architecture, training process, data provenance, compute budget, evaluation design, benchmark contamination, safety testing, known failure modes, energy consumption, model updates, and human interventions. It might include access to weights, code, system prompts, or controlled external testing. It might include independent replication by groups without financial ties to the company. It might include a public record of incidents, corrections, and withdrawn claims.

A conventional paper could contain some of this information. Frequently, however, it does not.

The scientific value of a disclosure therefore depends less on its format than on what it permits others to know and do. Can an independent investigator reconstruct the claim? Can the investigator test it under altered conditions? Can hidden assumptions be identified? Can failures be documented? Can the company's public language be compared with the underlying evidence? Can the result be challenged without requiring privileged access?

These are questions of epistemic accountability. Publication is one possible answer, but it is not the definition of the problem.

3. The Bibliometric Illusion

Bibliometric measurements create an appearance of precision. They convert a complicated institutional landscape into counts, percentages, rankings, and citation distributions. The resulting numbers feel objective because they can be reproduced, plotted, and compared. Yet the apparent neutrality of the method can conceal the conceptual choices on which the measurement depends.

To count the scientific output of AI startups, researchers must first decide which companies qualify as AI companies, which firms qualify as unicorns, which years matter, which databases are sufficiently comprehensive, which publications belong to which companies, and which authorship positions indicate a substantial contribution. Each choice determines what becomes visible.

A private company valued above $1 billion is a financial object, not a scientific category. The label “AI unicorn” may include frontier-model laboratories, autonomous-vehicle developers, enterprise-software vendors, medical technology firms, computer-vision companies, data platforms, defense contractors, and businesses that use AI primarily as a component of a larger commercial product. These organizations do not share a common scientific mission.

Some present themselves as research laboratories. Others sell services. Some make claims about scientific discovery, intelligence, medicine, or public safety. Others merely offer automation tools. Treating them as a single scientific population creates a category whose apparent coherence derives largely from investor valuation.

The valuation threshold itself contributes nothing to scientific interpretation. A company does not become more obligated to publish because investors estimate its worth at $1.1 billion rather than $900 million. Nor does valuation indicate the size of its research workforce, the originality of its technology, or the public consequences of its claims.

The category is useful for producing an arresting headline because unicorns symbolize wealth, ambition, and hype. It is less useful for determining what scientific standards ought to apply.

The denominator used in the article creates a similar illusion. The startups reportedly produced only about one in every thousand AI papers published during the measured year. The comparison sounds devastating. Yet the denominator includes a vast global literature generated by universities, established corporations, hospitals, government laboratories, engineering departments, and researchers applying AI across hundreds of fields.

There is no obvious reason that a few hundred private startups should produce a large fraction of the entire world’s literature involving artificial intelligence. The small proportion is therefore

not necessarily evidence of extraordinary silence. It is partly a predictable result of comparing a narrow financial category with an immense and heterogeneous scholarly corpus.

The statistic is rhetorically effective because readers are invited to interpret one in a thousand as a measure of absence. But the ratio does not establish that the companies contribute only one-thousandth of relevant innovation, technical capability, research investment, or scientific influence. It measures only their share of a database constructed under specific inclusion rules.

Bibliometrics can describe the literature. They cannot, without additional evidence, describe the totality of knowledge production occurring beyond it.

4. From Publication Scarcity to Moral Suspicion

The article does not merely report that many AI startups publish few papers. It surrounds that finding with a cluster of morally charged associations: Theranos, fraud, secrecy, safety, energy use, social harm, and the possibility that advanced AI systems may threaten society. These concerns are not invented. Each is legitimate in its own domain. The rhetorical problem is that they are assembled around a study that does not directly measure any of them.

The invocation of Theranos is especially powerful. Theranos has become a cultural symbol of technological deception: a company that made extraordinary scientific claims while withholding the evidence required to test them. Mentioning it immediately raises the stakes. Sparse publication no longer appears merely as a difference in institutional practice. It begins to look like a warning sign of possible fraud.

There are circumstances in which that inference is justified. A company claiming to diagnose disease, discover drugs, estimate mortality, or replace established scientific procedures should be expected to provide evidence proportionate to the consequences of its claims. In medicine, secrecy can directly endanger patients. In such cases, publication, regulatory review, data disclosure, and independent validation are not ceremonial expectations. They are safeguards.

But the analogy becomes misleading when applied indiscriminately across the entire category of AI unicorns. An enterprise software company, a coding assistant, a logistics platform, and a medical diagnostic firm do not create the same epistemic or public risks. They should not be judged by identical disclosure standards.

The article benefits from moving rapidly across these distinctions. It begins with paper counts, passes through Theranos, and arrives at questions of safety and social impact. The reader is encouraged to experience these topics as parts of a single problem. Yet publication scarcity may have different meanings in different settings.

For a medical company, the absence of public validation may be intolerable.

For a frontier AI laboratory, the absence of meaningful safety disclosure may be dangerous.

For a commercial software company, the absence of journal articles may be unremarkable.

For an open-source project, a repository, documentation, and public issue history may be more revealing than a conventional paper.

The moral evaluation should follow the claim and the risk, not the company's valuation category.

A stronger analysis would therefore begin by identifying what each company claims to have achieved. It would then ask what evidence is publicly available, what independent testing has occurred, what harms could result if the claim is false, and what form of disclosure would be appropriate. The publication count would become one component of a broader accountability assessment rather than the primary moral indicator.

The article instead treats publication scarcity as a generalized sign of epistemic deficiency. That move is rhetorically efficient because it allows many anxieties to accumulate around a single number. It is analytically weak because the number cannot bear that weight.

5. Institutional Science as an Interested Party

The article presents institutional scientists mainly as evaluators of corporate behavior. Their own interests remain largely invisible.

This omission matters because the dispute is not occurring between a neutral scientific community and an external commercial sector. Universities, journals, professional societies, funding agencies, and academic laboratories are themselves institutions competing for money, attention, authority, personnel, and influence.

Academic science has much to lose from the migration of discovery into private organizations. Corporations can offer higher salaries, larger computing budgets, faster development cycles, proprietary data, and direct control over infrastructure. They can recruit researchers before academic departments can hire them, absorb methods before journals can publish them, and deploy systems before universities can study them.

This produces a structural anxiety. The university retains immense symbolic authority, but it increasingly lacks material control over the most expensive forms of AI development. The question of publication therefore becomes a proxy for a larger struggle over who possesses legitimate knowledge.

The academic response is understandable. Scientific knowledge should not disappear behind corporate walls. Claims affecting medicine, public policy, labor, education, communication, and national security should not be insulated from scrutiny merely because they were produced privately. But the defense of public knowledge can easily become entangled with the defense of academic status.

The article's emotional appeal depends on that entanglement. It tells institutional scientists that the companies may possess wealth and power, but they do not fully possess science because they do not participate sufficiently in the publication system. Their work may be commercially impressive, but it has not been properly entered into the recognized record of knowledge.

This is comforting because it preserves a hierarchy. The corporation can build, but the academy decides what counts. The corporation can scale, but the journal certifies. The corporation can attract capital, but the citation system confers intellectual legitimacy.

The difficulty is that academic publication itself is not outside the prestige economy. It is one of its central mechanisms.

Researchers compete for authorship position, journal placement, citation counts, grants, tenure, prizes, and institutional rankings. Universities advertise publication volume. Journals market impact factors. Funding agencies reward visible output. Citation systems amplify already prominent researchers and institutions. Peer review often reinforces disciplinary expectations and prestige hierarchies.

None of this means that academic science is fraudulent. It means that it is socially organized and institutionally interested. Its norms often serve both knowledge and career.

A serious critique of corporate AI should acknowledge this duality. Otherwise, publication is idealized as a pure epistemic practice while corporate secrecy is interpreted as a uniquely commercial corruption. The reality is more complicated. Both systems contain mixtures of inquiry, prestige, secrecy, cooperation, competition, and self-protection.

Universities conceal data until publication. Researchers delay disclosure to preserve priority. Journals place knowledge behind paywalls. Laboratories withhold code. Institutions issue press releases before independent replication. Negative results vanish because they offer little career reward. Replication work is undervalued because novelty carries more prestige.

Corporate laboratories have their own distortions: trade secrecy, investor pressure, product timelines, competitive withholding, legal risk, marketing incentives, and resistance to external evaluation. But the contrast is not between corrupt commerce and pure science. It is between two imperfect institutional systems with different incentives and different forms of opacity.

The proper question is therefore not which institution deserves automatic trust. It is which arrangements make consequential claims most open to correction.

6. The Self-Flattering Function of the Headline

Clickbait is not defined only by exaggeration. It often works by giving readers a conclusion they are already prepared to enjoy.

The headline about AI startups "barely publishing" offers institutional scientists a satisfying reversal. The companies may dominate news coverage, investment, product development, and public imagination, but they are depicted as deficient according to the academy's own measure of seriousness.

This rhetorical structure is especially effective because it converts a material disadvantage into a moral advantage.

Universities may lack the compute resources of frontier laboratories.

They may struggle to retain researchers.

They may depend on corporate APIs and platforms.

They may be unable to reproduce the largest models.

Yet they can still say: the companies are not doing science properly.

The publication metric thus becomes a compensatory instrument. It restores status where material power has been lost.

The article strengthens this effect by using language such as "one of science's most fundamental practices." That phrase does more than describe publication. It elevates it into a defining obligation. A company that does not publish is not simply choosing a different communication channel. It is failing to participate in science at a fundamental level.

But science has never been reducible to one publication system.

Scientists have communicated through letters, notebooks, monographs, lectures, instrument demonstrations, laboratory reports, field records, patents, engineering drawings, code, databases, and informal correspondence. Different fields have developed different norms. In mathematics, a proof may be decisive. In experimental biology, materials and methods may matter most. In computer science, code and benchmark performance may be central. In engineering, a working artifact may reveal capacities that a paper only describes.

The scientific value of an output depends on the kind of claim being made and the kind of scrutiny that claim requires.

The headline suppresses this diversity because a narrower formulation would be less dramatic. “Many AI unicorns publish few papers in indexed academic venues” is accurate but not emotionally powerful. It does not imply hypocrisy, secrecy, or scientific illegitimacy. It does not invite the reader to condemn a sector.

The stronger headline succeeds by allowing “publishing” to stand for “doing accountable science.” Once that substitution is accepted, the companies appear guilty before the details are examined.

This is a familiar pattern in institutional journalism. A metric is selected, a population is ranked, a dramatic disparity is found, and the disparity is framed as a moral revelation. The number gives the story authority. The framing gives it emotional direction.

The reader is not merely informed that AI startups publish little. The reader is invited to feel that the startups have been exposed.

That is the clickbait function.

7. The Problem of Selective Skepticism

The article applies skepticism asymmetrically.

Corporate claims are treated as suspect because they are not sufficiently documented through scientific literature. Academic publication, however, is treated as the remedy rather than as another object requiring evaluation.

This asymmetry obscures a central fact: publication can produce the appearance of accountability without its substance.

A paper may contain enough information to satisfy reviewers but not enough to reproduce the work. It may rely on private data. It may report benchmark results without disclosing prompt selection, excluded trials, failed runs, or manual interventions. It may describe a model architecture while withholding the training process that determines its actual behavior. It may use statistical significance to conceal weak practical effects. It may present a carefully selected demonstration as though it represented general performance.

Once published, however, the work acquires the institutional symbols of legitimacy. It receives a citation. It enters a database. It becomes countable.

This creates a danger: the literature can be mistaken for the underlying reality.

The same problem appears in academic AI. Large numbers of papers may be produced around benchmarks that are saturated, contaminated, weakly connected to real-world performance, or easy to optimize without meaningful conceptual progress. Publication volume can rise while understanding remains stagnant.

A company that publishes ten benchmark papers is not necessarily more transparent than a company that releases a working model, extensive code, and a detailed evaluation repository. A university laboratory producing many papers is not necessarily more scientifically accountable than an engineering group that exposes its system to broad external testing.

The proper standard must therefore be format-neutral.

A paper should not receive credit merely because it is a paper. A repository should not receive credit merely because it is public. A model should not be called open merely because weights are downloadable. Each form of disclosure should be judged by what it allows outsiders to verify.

The central test is practical:

Can another investigator determine what was done?

Can the investigator identify what remains hidden?

Can the principal claims be challenged?

Can failures be found and reported?

Can the result be reproduced or meaningfully approximated?

Can conflicts between marketing language and technical evidence be exposed?

These questions apply equally to companies and universities.

Selective skepticism weakens the article because it directs scrutiny outward while leaving the publishing system itself largely unexamined. The result is not a defense of science in general. It is a defense of one institutional form of science.

8. Corporate Secrecy Is Real

Rejecting the article's framing should not become a defense of corporate opacity.

Many AI companies do conceal information that should be public. They release systems without adequate documentation, make claims that cannot be independently tested, change models

without clear notice, restrict access to evaluation details, and use the language of science in marketing materials while withholding the evidence that would make those claims scientific.

The closed nature of frontier AI raises serious problems.

External researchers may not know what data were used, what labor conditions supported data preparation, what copyrighted materials were incorporated, what safety failures were discovered internally, what interventions were required to produce benchmark results, or how much energy and water were consumed.

Model behavior may change through silent updates. Evaluation access may depend on company permission. Researchers who discover harmful behavior may face legal or contractual constraints. Public demonstrations may be curated. Safety reports may emphasize favorable findings while omitting adverse internal evidence.

These are genuine accountability failures.

But counting papers is a poor way to measure them.

A company could increase its publication total without disclosing the information that matters most. It could publish peripheral research while keeping its central systems secret. It could produce papers authored by a small internal research group while the main product organization remains opaque. It could publish safety work that cannot be independently reproduced because access to the relevant model is restricted.

Publication may even function as reputation management. A company can point to a respected scientific record while avoiding scrutiny of its most consequential operations.

The relevant distinction is therefore not publication versus secrecy. A company may publish and remain secretive. Another may avoid journals but release enough technical material to support substantial outside investigation.

The deeper issue is whether disclosure is proportionate to power.

A small software company selling a low-risk tool does not require the same transparency as a company deploying systems into medicine, education, military planning, financial markets, public administration, or mass communication. The greater the potential harm, the stronger the obligation to document, test, and expose the system to independent review.

This obligation should not depend on whether the company prefers journals, blogs, repositories, regulatory filings, or technical reports. It should depend on whether the evidence is adequate.

The article reaches toward this concern but cannot fully articulate it because its central metric remains publication count. It identifies a real anxiety with the wrong instrument.

9. Science as a Verb, Publication as a Credential

The confusion can be clarified by distinguishing science as an activity from publication as a credential.

Science as an activity includes observation, hypothesis formation, measurement, experimentation, model construction, criticism, replication, correction, and revision. It is a process through which claims are exposed to possible failure.

Publication is one way of recording and presenting that process. It becomes scientific when it enables examination and correction. It becomes ceremonial when it mainly signals institutional membership.

This distinction explains why publication counts are insufficient. A paper may carry the credential of science without embodying much of its corrective process. A technical artifact may embody that process without appearing in a journal.

A working software system can be tested by thousands of users. A public codebase can be inspected line by line. A model with accessible weights can be evaluated under conditions its creators did not anticipate. An open dataset can be reanalyzed. A detailed failure log can be more scientifically valuable than a polished publication that reports only successful outcomes.

At the same time, public availability does not guarantee scientific value. Code may be undocumented. Weights may be released without training data. Benchmarks may be cherry-picked. Users may test a model without understanding hidden system controls. A product may be observable but not reproducible.

The distinction is not between old and new formats. It is between outputs that support correction and outputs that resist it.

Science is strongest when claims are made vulnerable.

A scientific claim should be vulnerable to contradictory evidence, independent testing, methodological criticism, replication failure, and revision. Institutions that preserve claims from such vulnerability are behaving unscientifically, whether they are corporations, universities, journals, or governments.

This gives a more demanding standard than publication.

The question is not whether a paper exists.

The question is whether the claim has been placed in a condition where knowledgeable outsiders can make it fail.

10. From Publication Metrics to Epistemic Audits

A more useful investigation would replace the publication count with an epistemic audit.

Such an audit would begin with the claims a company actually makes. Does it claim that a system discovers drugs, writes reliable software, performs scientific reasoning, reduces energy use, improves diagnosis, surpasses experts, or approaches general intelligence? Each claim creates a corresponding burden of evidence.

The audit would then ask what an informed outsider needs in order to evaluate that burden.

A claim about medical performance may require representative patient data, predefined endpoints, comparison with accepted standards, subgroup analysis, external validation, adverse-event reporting, and regulatory scrutiny. A claim about scientific discovery may require disclosure of the search procedure, candidate-selection process, failed predictions, laboratory confirmation, and the human knowledge supplied before the system produced its apparent result. A claim about software performance may require public test suites, security audits, failure distributions, and comparisons under realistic conditions rather than staged demonstrations.

The required evidence changes with the claim. A single universal metric cannot capture it.

An epistemic audit would examine at least five broad questions.

First, what does the company say the system can do?

Second, what evidence supports that claim?

Third, which parts of the evidence can outsiders inspect?

Fourth, what material information remains inaccessible?

Fifth, what mechanisms exist for discovering, reporting, and correcting error?

This approach avoids the false choice between journal publication and corporate secrecy. A paper may contribute to the audit, but so may code, data, weights, independent evaluations, incident reports, regulatory submissions, and documented replications. The format matters less than the evidentiary function.

The audit must also distinguish disclosure from access. A company may disclose that it conducted safety testing without releasing the tests. It may describe training data without identifying their provenance. It may provide an application programming interface while preventing researchers from examining the underlying system. It may publish benchmark results while withholding the prompts, exclusions, model version, and selection process used to produce them.

These are partial disclosures. They may be useful, but they should not be mistaken for full verification.

The same scrutiny should be applied to academic claims. A university paper based on unavailable data, inaccessible code, undisclosed preprocessing, or proprietary models should not be treated as fully accountable merely because it passed peer review. Institutional affiliation does not repair missing evidence.

An epistemic audit would therefore create a common standard across organizations. It would neither privilege corporations because they build functioning products nor privilege universities because they publish conventional papers. It would ask what can actually be known.

11. The Politics of Epistemic Legitimacy

The dispute over publication is ultimately political because the authority to define science carries social power.

Institutions recognized as scientific can influence public policy, receive government funding, shape education, advise courts, define professional standards, and determine which claims deserve public trust. To call an activity scientific is not merely to classify it. It is to grant it a privileged position in the organization of society.

The rise of commercial AI destabilizes this arrangement. Private laboratories increasingly produce systems that affect knowledge itself: how information is retrieved, summarized, ranked, translated, generated, and interpreted. These companies do not simply make tools used by scientists. They increasingly mediate the public's access to scientific and cultural reality.

Their power therefore exceeds the conventional boundaries of industrial research.

At the same time, universities and journals no longer possess an uncontested claim to epistemic authority. Researchers publish into an environment saturated with preprints, repositories, independent analysts, open-source communities, social media, private laboratories, and machine-generated synthesis. The journal article remains important, but it is no longer the sole practical gateway through which new knowledge enters public circulation.

Institutional science can respond to this transformation in two ways.

It can defend its inherited procedures as the defining form of legitimate knowledge.

Or it can defend the deeper principles those procedures were intended to serve: transparency, criticism, replication, correction, and public intelligibility.

The first response protects institutional status. The second protects science.

The difference becomes visible in the treatment of outsiders. A publication-centered institution asks whether a claim has entered the approved literature. A science-centered institution asks whether the claim has survived serious attempts to test it.

These questions often overlap, but not always.

The first can exclude unconventional contributors whose work is public, testable, and valuable but lacks institutional affiliation. It can also admit weak claims that have successfully navigated professional procedures. The second is harder to administer because it cannot be reduced easily to journal names, author positions, or citation counts. It requires direct engagement with evidence.

This is one reason bibliometrics are institutionally attractive. They make evaluation scalable. Administrators can count papers without understanding them. Funding agencies can compare citation indices without reproducing experiments. Journal prestige can substitute for direct judgment. The metric economizes attention.

Corporate AI intensifies the limitations of this system because its most consequential outputs may be systems rather than papers. A model can alter millions of decisions before the academic literature has adequately characterized it. A product can change repeatedly while the published description remains fixed. A safety paper can become obsolete after a model update. A benchmark result can be surpassed, invalidated, or revealed as contaminated within weeks.

The static publication record is poorly matched to rapidly changing systems.

This does not make scientific literature obsolete. It means that accountability must become continuous. Models should be versioned. Evaluations should be repeated. incidents should be recorded. material changes should be disclosed. Independent access should be preserved. Claims should be revisited as systems evolve.

The scientific object is no longer always a stable experiment reported once. It may be a changing technical service embedded in society.

The institutions of science must adapt to that fact rather than merely demanding that corporations produce more papers.

12. Beyond the Academic–Corporate Binary

The article’s framing depends on a simplified opposition between commercial AI and institutional science. In reality, the two are deeply interdependent.

Academic researchers move into corporations and sometimes return. Universities receive industrial funding. Companies sponsor conferences, provide grants, donate computing resources, release tools, and recruit from publicly supported laboratories. Academic papers rely on corporate hardware, software, cloud services, and proprietary models. Corporate research draws on decades of publicly funded mathematics, computer science, physics, linguistics, psychology, and engineering.

The boundary is porous.

This makes moral storytelling difficult. It is convenient to portray the university as the home of openness and the company as the home of secrecy, but both institutions participate in mixed economies of disclosure.

A university may announce an open-science commitment while restricting access through journal subscriptions, expensive tuition, proprietary databases, patent offices, and exclusive commercial partnerships. A company may protect its central technology while releasing code, weights, technical documentation, and tools that enable research across the world.

Neither institutional type has a monopoly on openness.

A broader ecology of knowledge has also emerged outside both sectors. Independent researchers, nonprofit organizations, open-source communities, citizen scientists, small laboratories, online collaborators, and unaffiliated scholars increasingly contribute to technical and scientific work. Their contributions may appear in repositories, videos, public notebooks, discussion forums, software releases, and informal preprints.

A publication-centered account risks making these contributors invisible.

The bias is structural. Bibliographic databases are designed to count recognized documents. They are less capable of recording distributed development, iterative correction, software maintenance, public debugging, visual demonstration, or knowledge that emerges through collaboration outside formal institutions.

This does not mean every repository is science or every independent claim deserves credibility. It means that scientific value must be assessed at the level of evidence and method rather than inferred from institutional location.

A mature account of AI knowledge production would therefore distinguish among several overlapping systems:

academic publication, corporate research, commercial engineering, open-source development, regulatory evaluation, independent auditing, and public experimentation.

Each system has strengths and weaknesses.

Academic publication can preserve claims and expose them to specialized criticism, but it is slow and vulnerable to prestige incentives.

Corporate engineering can build systems at scale, but it is vulnerable to secrecy and marketing pressure.

Open-source development can enable broad inspection, but it may lack resources, documentation, or stable governance.

Regulation can compel disclosure, but it may become bureaucratic, politically captured, or technologically outdated.

Independent auditing can identify failures, but auditors require access, funding, and legal protection.

No single system is sufficient.

The appropriate response is not to restore the journal as the unquestioned center of scientific legitimacy. It is to build a plural accountability structure in which different forms of evidence reinforce and challenge one another.

13. Conclusion: The Right Indictment

The article begins with a valid observation and converts it into a comforting morality play.

Many highly valued AI startups publish few conventional scientific papers. Their visible participation in indexed literature is sparse, concentrated, and often dependent on a small number of authors. This deserves attention, especially when companies make claims about science, medicine, intelligence, safety, or social transformation.

But publication scarcity is not the same as absence of research. It is not the same as absence of documentation. It is not the same as irreproducibility. And it is not, by itself, evidence of fraud, danger, or scientific emptiness.

The article becomes clickbait when it encourages those meanings to collapse into one another.

Its emotional reward is equally important. It allows institutional scientists to interpret corporate power as epistemic inferiority. The companies may possess capital, infrastructure, personnel, and public influence, but the academy can still claim the superior moral position because it publishes.

That reassurance is understandable, but it is not sufficient.

Institutional science should not answer the rise of commercial AI by treating its own publication system as the definition of knowledge. It should answer by demanding stronger forms of public accountability from every institution, including itself.

The proper indictment of AI companies is not that they fail to imitate universities. It is that many exercise extraordinary epistemic and social power without making their consequential claims sufficiently vulnerable to independent challenge.

The proper indictment of institutional science is not that it publishes too much. It is that it sometimes confuses publication with truth, citation with importance, peer review with verification, and professional recognition with public accountability.

Both systems can generate knowledge.

Both systems can conceal error.

Both systems can use the language of science to protect prestige and power.

The decisive standard must therefore lie deeper than institutional form.

A claim deserves scientific standing to the extent that its evidence can be inspected, its methods can be understood, its limitations can be identified, its failures can be reported, and its conclusions can be revised when reality resists them.

Science is not secured by the existence of a paper.

It is secured by the organized possibility of correction.

References

Ahmed, N., Wahed, M., and Thompson, N. C. “The Growing Influence of Industry in AI Research.” *Science*, 2023.

Ioannidis, J. P. A. “Why Most Published Research Findings Are False.” *PLoS Medicine*, 2005.

Ioannidis, J. P. A. “How to Make More Published Research True.” *PLoS Medicine*, 2014.

Merton, R. K. “The Normative Structure of Science.” In *The Sociology of Science: Theoretical and Empirical Investigations*. University of Chicago Press, 1973.

National Academies of Sciences, Engineering, and Medicine. *Reproducibility and Replicability in Science*. National Academies Press, 2019.

Nosek, B. A., et al. “Promoting an Open Research Culture.” *Science*, 2015.

Open Science Collaboration. “Estimating the Reproducibility of Psychological Science.” *Science*, 2015.

Resnik, D. B. *The Ethics of Science: An Introduction*. Routledge, 1998.

Science. “AI’s Top Startups Are Barely Publishing Their Research.” July 29, 2026.

Ziman, J. *Real Science: What It Is, and What It Means*. Cambridge University Press, 2000.